CHAPTER

1

DEFINING A SERVER

After reading this chapter and completing the exercises, you will be able to:

- Differentiate between peer-to-peer and client-server networking models
- Identify server functions and benefits
- Identify characteristics that distinguish server hardware from client hardware
- Identify three main types of servers

The server is the single most critical element of most networks. This fact becomes very evident when a server fails. If a user's computer fails, it affects only that user. However, if a server fails, it can affect hundreds or thousands of users and disrupt an organization's business operations. In server planning, you must be able to determine if the network needs a server (or an additional server) and what server capabilities are required. What kind of hardware will meet the anticipated needs? How will you justify the initial and maintenance expense of server equipment? In order to make these decisions, you should know what to look for in a server. This chapter starts with the most basic network model—a simple network without a server—and defines the benefits of adding a server. You will also learn the characteristics that distinguish a user's computer from a server, and what distinguishes various server types from one another.

Does Your Network Need a Server?

Before you delve into planning and installing a server, consider if you really need a server. This might seem like an odd statement to place in the first chapter of a book on servers, but it is a legitimate one. If your organization is small (perhaps 12 or fewer workgroup computers), users only occasionally share printers, files, or applications, and security is not a major concern, then it is possible to network all the computers together without a server. For example, users in a small, family-run construction company probably already trust one another and only need to print invoices or checks to a shared printer from time to time. The company does not require a database, and the construction managers create job estimates on their own laptop PCs, which only occasionally connect to the network. In this case, it is of little use to implement a server.

Conversely, most organizations (even small ones) can benefit from adding one or more servers to the network. Some of the benefits include enhanced security, improved performance, centralized file storage, centralized administration, and a central location from which to run applications. For example, the owner of a small real-estate office installs a high-speed Internet connection so that real-estate agents can research land values and receive email from clients. Also, the owner wants each agent to enter sales records into a company database. A server could include special software to protect office users and the company database from malicious Internet users and provide a central location to store the company database.

After determining the need for a server, you must also determine "how much" server you need. Will the server support dozens, hundreds, or thousands of users? Do you require extremely high performance? Some servers offer features and performance comparable to common desktop personal computers, some are extremely powerful and cost hundreds of thousands of dollars, and others fall somewhere in the middle.

Comparing Peer-to-Peer and Client-Server Models

A **network** is a collection of two or more computers connected with wired transmission media such as network cable or with wireless radio or infrared signals, and it usually includes other devices such as printers. A **network device** is any device connected to the network for purposes of communicating with other network devices. (A network device is also known as a **host** in most networks.) Cumulatively, the network devices form what is known as a **local area network (LAN)**—a collection of computers in close proximity to one another on a single network (see Figure 1-1). A LAN includes a single cable to which all computers connect, or more commonly, each device has its own network cable, each of which converges to a central **hub** or **switch**. You can use a network cable to join multiple hubs together, but the entire collection of hubs still constitutes a single network. A LAN can be one of two basic networking models: peer-to-peer or client-server. Expanding upon a LAN, a **wide area network (WAN)** involves multiple, geographically distant LANs connected to one another across relatively great distances. For example, an organization's LAN in Seattle connected to the same organization's LAN in Phoenix constitutes a WAN.

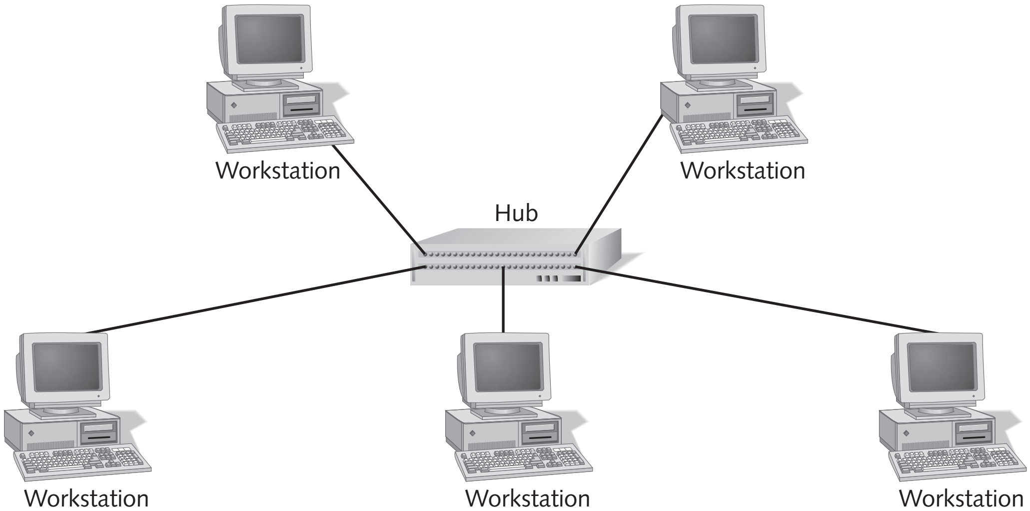

Figure 1-1 A simple LAN connected to a hub

Peer-to-Peer

A network does not necessarily require a server. A **peer-to-peer network** is a network of computers with no logon server to verify the identity of users. (Although a peer-to-peer network might include a file server that stores files for the users, this is seldom the case.) This model is called a peer-to-peer network because each network device has an equal (peer) level of authority. Computers in a peer-to-peer network are usually common desktop computers, otherwise known as **workstations**, and are generally equipped with only enough hardware to service the needs of a single user. Most peer-to-peer networks use the file and printer sharing capability of Microsoft Windows 95/98/ME/NT/2000 to share network resources. A **network resource** (sometimes called only a "resource") is an object that users can access across the network. Common examples of network resources include printers, files, and folders (see Figure 1-2).

Although the file and printer sharing capability of a peer-to-peer network can service a small network, it can also limit network growth. For example, Microsoft Windows 95, 98, NT Workstation, and 2000 Professional limit inbound concurrent network connections to 10. This means that if 10 users are already accessing a folder, an additional attempt to access the folder will fail because the limit of connections has already been reached. This limitation restricts the size and growth capacity of peer-to-peer networks.

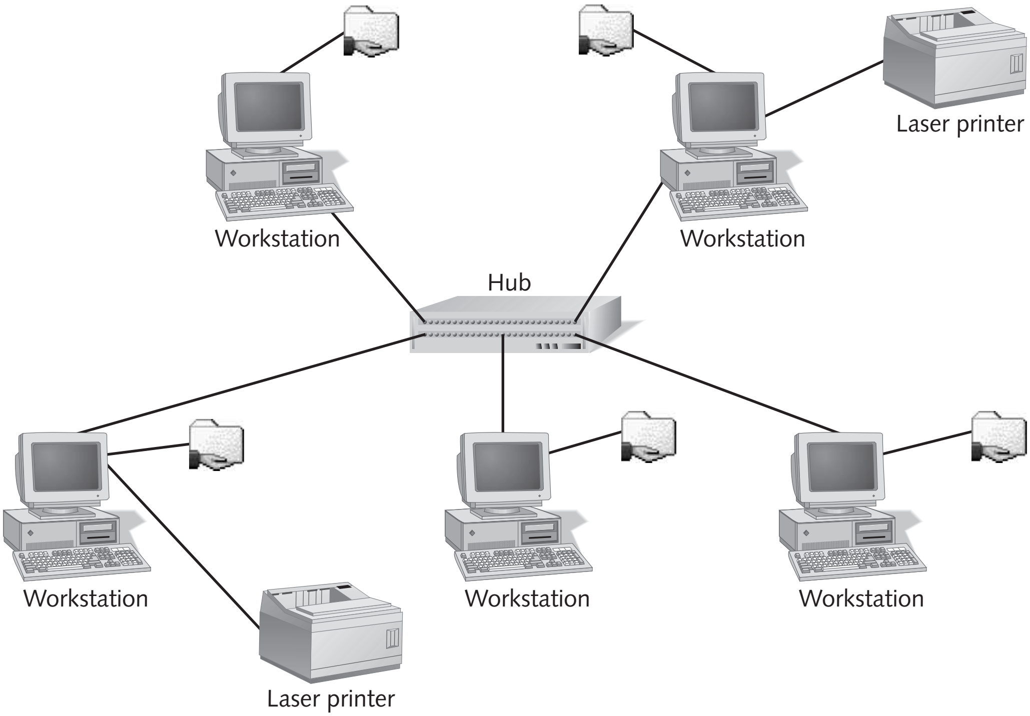

Figure 1-2 Users on the LAN can share folders, files, and printers

Peer-to-peer networks function best when individual users seldom access one another's files. If user A accesses a file on user B's workstation, then saves the file to user A's local hard disk, it causes confusion when attempting to locate the most up-to-date version of a file.

Peer-to-peer networks also lack file security. Windows 95/98/ME shares folders (and hence the files contained in the folders) with password protection. A password can be easily compromised and, at best, is only as secure as the operating system. Most desktop client operating systems are designed with minimal security. For example, any passerby can boot up a Windows desktop client and, whether logging on to the network or not, access local resources on the entire local hard disk.

Windows NT or 2000 workstations in a peer-to-peer configuration, on the other hand, require you to create a user account that has a username and password in order to log on locally. Access to folders or individual files can be specified for designated users. If a user attempting to access your shared files across the network does not have valid credentials for the shared files, he or she is denied access. Although Windows NT or 2000 is inherently more secure than Windows 95/98/ME password security, a peer-to-peer network requires you to store user accounts on each workstation.

For example, suppose you, Shelly, and Karl all have Windows NT 4.0 workstations in a peer-to-peer network. You want to share a folder with Shelly, and Karl also wants to share a folder with Shelly. Both you and Karl must create a user account for Shelly on your respective workstations, and Shelly's username and password must be exactly the same. If Karl changes the password for Shelly's user account on his machine and you don't, then Shelly's credentials will be inconsistent and she will not be able to access resources where her password is different than the one she used to log on (see Figure 1-3). The root of this problem is that there is no centralized location in which to store user accounts, a significant weakness of the peer-to-peer network. Managing users and resources without centralized user account management can quickly become a time-consuming nightmare.

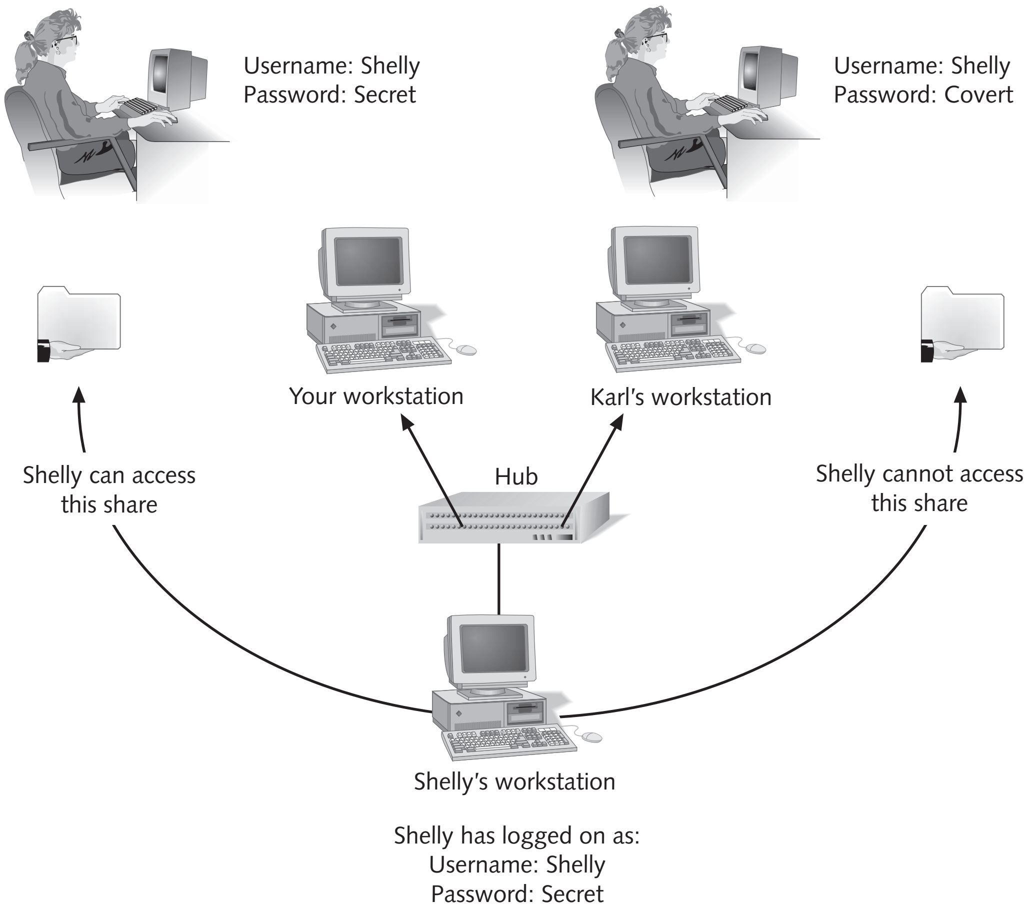

Figure 1-3 Inconsistent user accounts can cause access problems

A peer-to-peer network is simple to configure, but it is limited in terms of expandability, features, services, and security. Many small networks are peer-to-peer networks because the

organizations they serve do not require (or cannot afford) the benefits of a server, or there is insufficient technical expertise to administer a server. A peer-to-peer network:

- Usually involves a small number of computers (around one dozen or less)
- Has limited growth potential
- Has decentralized file management, user accounts, and overall management
- Offers minimal security
- Is simple to configure
- Is typically the least expensive option

Client-Server

The **client-server** networking model has most of the benefits of the peer-to-peer network model and potentially none of its weaknesses. The client-server network begins with a LAN and one or more servers, but it can also encompass a more complicated network configuration such as a WAN.

A **server** usually possesses more processing power, RAM, and hard disk capacity than workstation computers on the LAN. The server also has a server **network operating system (NOS)** such as Microsoft Windows NT or 2000, Linux, IBM OS/2, or Novell NetWare. A server running a NOS provides file and printer sharing, centralized file storage, administration, security, services, and significantly more stability than desktop operating systems.

At a basic level, simply installing a server into an existing peer-to-peer network changes the network model to a client-server network (see Figure 1-4). Workstations in the client-server network request and receive services (access to files, printers, or applications, for example) from the server—hence the term **client**. The server serves the client—hence the term server.

Although users at the workstations can save files on their local hard disks, it may be more practical to save files on the server so that all users have a central location to access files. A server in the role of storing files on behalf of network clients is commonly referred to as a **file server**. A file server also offers other advantages, including:

- *Version control:* Using a file server helps to avoid the situation in which differing versions of a file are saved on various workstations.
- *Backup:* Network administrators can back up data files centrally on the file server as opposed to hunting for and individually backing up the data on each user's workstation.
- *Security:* Network administrators can centrally audit file resources to see if unauthorized persons have tried to gain access to them. Also, administrators can configure security settings on the files instead of users, improving the chances that sensitive files will be well protected.

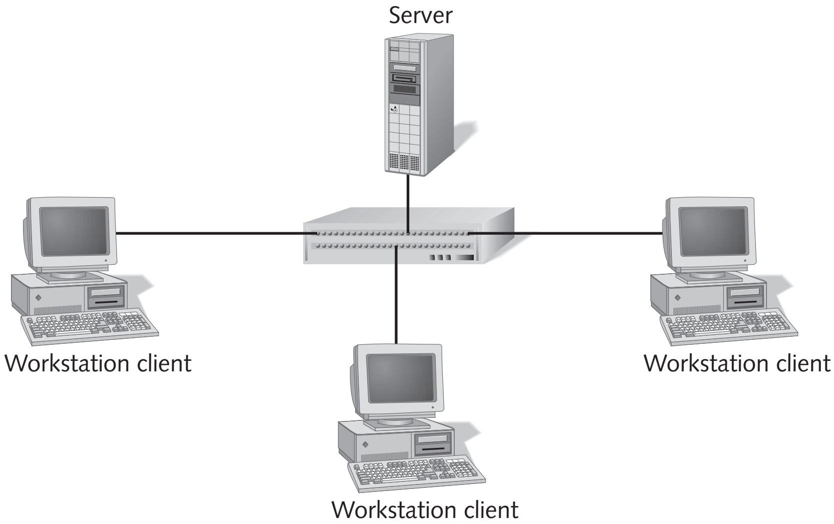

Figure 1-4 A client-server network

- *Availability:* The file server is usually placed in a well-connected location to increase the availability of file resources. Because of the stability of the NOS and the use of redundancy, the server is less likely to be subject to outages or downtime. (**Redundancy** is the ability to continue providing service when something fails. For example, if a hard disk fails, a redundant hard disk can continue to store and serve files.) In addition, the server remains on 24/7, so it is always available to clients, as opposed to a computer in a peer-to-peer network in which resources are only available when the user turns on the computer.
- *Integrity:* NOS file systems protect the integrity of their files by using special logs and error correction that can repair a damaged file on the fly or issue an alert to an administrator.

Because a NOS is capable of hundreds or thousands of simultaneous connections, servers can significantly increase the size and growth potential of a network. Some servers can even handle millions of simultaneous connections provided the server hardware and network bandwidth can match the demand.

A server in the client-server network model:

- Possesses more processing power, RAM, and hard disk capacity than typical workstations
- Uses a NOS such as Microsoft Windows NT, Windows 2000, Linux, IBM OS/2, or Novell NetWare
- Provides a central file storage location
- Is capable of many more simultaneous connections than a workstation

- Offers security features such as logon authentication
- Provides centralized administration

All these benefits are the result of server implementation in the network. In addition to these features, servers offer several other functions and benefits.

SERVER FUNCTIONS AND BENEFITS

The server's impact on the network varies depending upon which of its many features you choose to implement. Even adding a single server to a peer-to-peer network adds benefits to the network in terms of services, security, performance, storage capabilities, access to applications, and centralized management. These features add up to more cost effectiveness, efficiency, and productivity.

Services

A **service** is a function of the NOS that provides various server functions and benefits to the network. For example, a company web server could host many services. When a user accesses the web site by typing a Uniform Resource Locator (URL) into a web browser, the Domain Name System (DNS) service translates the URL into a unique number that identifies the server that contains the web pages. When the user views the web page, he or she might want to leave a message for someone at the site by clicking a link that sends an email message. The email functionality requires an email service. If the user purchases products from the web site and wants to view his or her account information, a service checks the user's username and password so that only that user can view the account. Yet another service could filter and monitor all traffic from the user's network to and from the Internet for security purposes. In most networks, these services are probably distributed among several servers for purposes of redundancy and performance.

Note: Specific services are addressed in more detail in later chapters.

Security

Protecting the network against unauthorized access is a significant challenge. Though you want to grant network access to users, you do not want to allow them free reign on the network and its resources. Conversely, too many restrictions will be counterproductive. The administrator must permit a level of access that is only as much as the user needs.

Authentication

Security starts with **authentication**, which verifies a person's identity based on their credentials entered at logon (usually a username and password). As an authenticator, the

server is an authority in the network, negating the concept of a peer-to-peer network in which all computers are equal. Once authenticated by the server, the user gains general access to the network for which the server is responsible. Thereafter, when the user attempts to access network resources, the authorized user account is compared against a list of users or groups with access to those resources. If the user account is on the list, the NOS grants access. This process of verifying that the user has the ability to access the resource is known as **authorization** (see Figure 1-5).

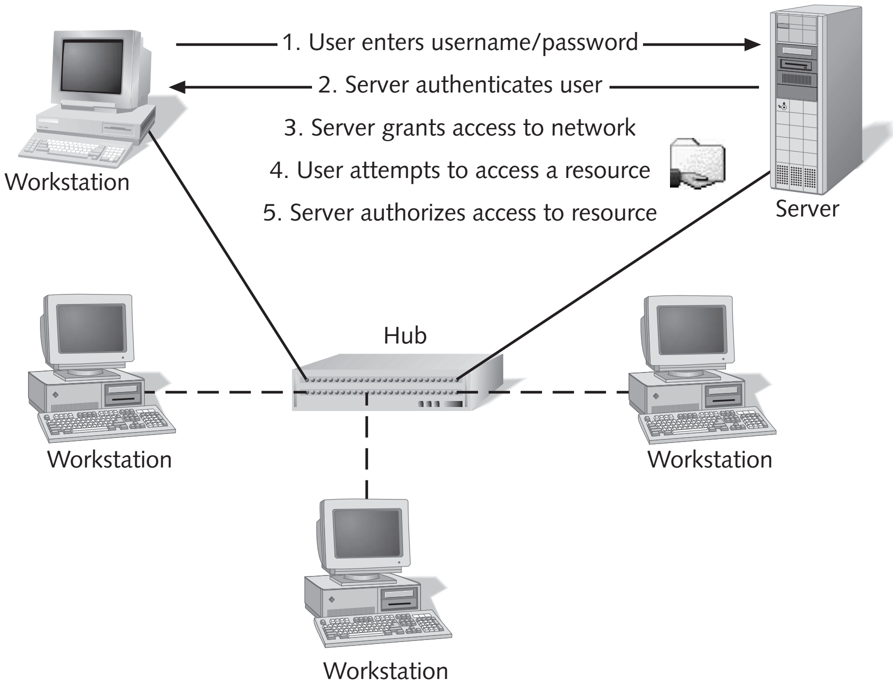

Figure 1-5 The server authenticates users and authorizes access to network resources

In a server environment, authentication better protects network resources than password protection in a peer-to-peer network. Each user account is stored on the server only, mitigating the need to create user accounts on each workstation. Recall that in a peer-to-peer environment, when a user wants to share a folder, he or she must protect it with a password or have sufficient administrative rights to create user accounts for those who want to access it. In a client-server environment, instead of using a password or creating user accounts, permission to access the folder can be granted by the user (or an administrator) to a list of users or groups of users on the server. The administrator creates the users and groups, and it is not necessary to grant administrative rights to individual users. Figure 1-6 shows a list of users (shown with a single head icon) and groups (shown with a double head icon) on a Windows 2000 Server computer.

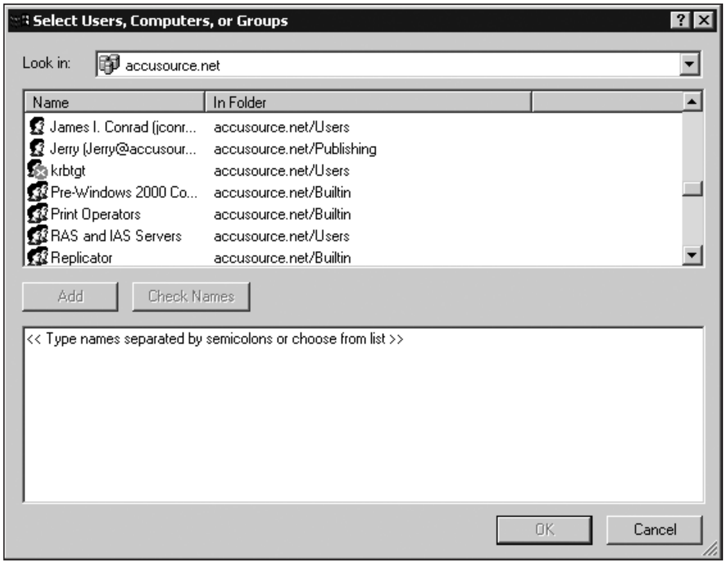

Figure 1-6 Granting access to users or groups

For added security, consider using emerging technologies such as smart cards to authenticate users (see Figure 1-7). A smart card includes an encrypted certificate, and the user typically inserts the card into a reader attached to the computer and enters a password or personal identification number (PIN). Because it requires both the physical card and a password or PIN, the smart card protects against unauthorized persons using stolen passwords. Similarly, a SecurID card synchronizes with the computer to generate a randomized pattern of six-digit numbers displayed in a timed sequence; you enter what is displayed along with a PIN as your password and you are granted access based on this "one-time" password. This "something you have and something you know" authentication is similar to smart cards, but insertion into a reader is not required.

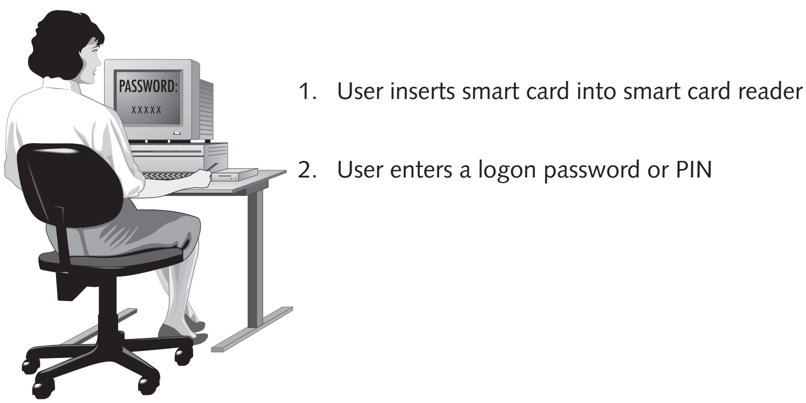

Figure 1-7 Smart card authentication provides a high level of security

Permissions

Authentication verifies the user's identity, and authorization compares that identity against a list of users or groups that are permitted to access resources such as files. Once a user accesses a resource, however, there is the further question of what that user can do with that resource. That is what permissions are about. **Permissions** specify the degree to which a user or group can access or alter a resource. For example, an administrator could set permissions on a folder for individual users or for groups. Who should be able to read but not change the folder contents? Who should be able to access the folder and change permissions? It could be extremely problematic if permission were given to the wrong users or groups to access sensitive information. Consider what would happen if everyone on a company network were able to read the Salaries.xls spreadsheet! To avoid this kind of trouble, administrators assign specific permissions to individuals or groups of individuals. For example, the administrator might assign permissions to the Human Resources department that allows the ability to read and change the Salaries.xls spreadsheet, assign read permissions to managers, and deny permissions to all other users. The NOS permits the administrator to assign these very specific permissions for resources, as shown in Figure 1-8.

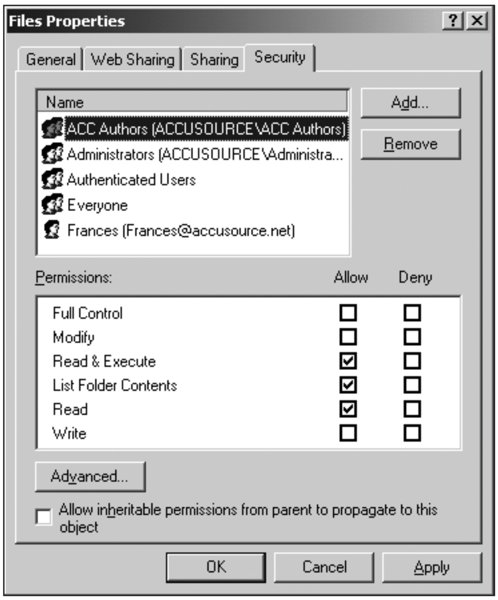

Figure 1-8 The NOS (Windows 2000 in this case) allows you to configure specific permissions

Of course, even on a peer-to-peer network, a user can set (perhaps unwise) permissions on a workstation resource. However, because a server centralizes resources, it makes

monitoring those resources more efficient. Compare monitoring permissions individually on 10 computers in a peer-to-peer network with monitoring permissions for tens or hundreds of workstations on a single server.

Security Boundaries

Closely associated with security, authentication, and permissions is the issue of security boundaries, which delineate where a user can "go" in the network. In a large enterprise, administrators do not want an employee in Los Angeles to be able to access a resource in Hong Kong only because they work for the same company and are connected to the same network. To control user access, administrators can use security boundaries such as the domain model of Windows NT 4.0. Users in the Los Angeles domain can access resources only in the Los Angeles domain. Similarly, users in the Hong Kong domain cannot access resources in the Los Angeles domain. However, the administrator of each respective domain can manually allow access between the domains using what is known as a **trust**. Windows 2000 also uses a domain system, except that it implements automatic trusts between domains—administrators do not have to manually create a trust.

Performance

Because a server typically has significantly more powerful processors, more hard disks, and greater storage capacity than a workstation, it can perform tasks on behalf of the workstation, improving overall performance and freeing the workstation to perform other tasks. For example, a special-effects technician creating a video clip can use the server to render the clip instead of using his or her own local workstation. High-end servers that provide better performance can also cost several times more than a low-end workstation. It is more financially feasible to purchase a $20,000 server to service 10 special-effects technicians than it is to spend $3,000 to upgrade each technician's individual workstation (for a total of $30,000). In this simple example, the administrator saves the organization $10,000 by using a server instead of upgrading individual workstations.

Processor

Three primary factors—clock speed, data bus, and cache—contribute to the effective speed of the processor (CPU).

- **Clock speed** is the number of cycles the processor can execute in a single second, measured in millions of cycles per second, or megahertz (MHz). Instructions executed by the processor require a certain number of cycles, so the more cycles the processor can handle per second, the faster it operates, or "thinks." Current processors are also **superscalar**—that is, they can execute more than one instruction in a single clock cycle.

Note

The processor operates at a speed in MHz that is a multiplier of the bus speed. For example, a 700 MHz processor multiplies a 100 MHz bus speed times seven.

- **Data bus** refers to the number of data bits that can pass into or out of the processor in a single cycle. Data bus width is typically 32 bits. Think of a 32-lane highway over which data travels. Some new processors offer 64-bit bandwidth, equivalent to a 64-lane data highway. Because of other internal engineering modifications, a 64-bit processor utilizes its resources more efficiently than a 32-bit processor and is, therefore, more than twice as fast. Data bus width and RAM have a direct relationship such that a 64-bit processor can also utilize 4,294,967,296 times more memory than a 32-bit processor (assuming that the 64-bit architecture extends throughout the entire system).
- **Cache** is memory that exists on or near the processor itself but is separate from main system memory (RAM). Processor cache stores recently accessed data from the hard disk or RAM. When the same information must be retrieved at a later time, the system can access instructions stored in the cache more quickly than the system can retrieve information from the hard disk or RAM. Also, the cache anticipates what the processor will request next and fetches it in advance, hoping that it has guessed correctly. (Intel claims that 90 percent of the time, the cache guesses correctly.)

Although clock speed, data bus, and cache are components of both servers and workstations, servers are typically more powerful in each respect. Also, servers often utilize **symmetric multiprocessing (SMP)**, which is the simultaneous use of multiple processors on the same server. SMP results in a corresponding increase in performance such that two processors are about twice as fast as one, four processors are about twice as fast as two, and so forth.

Note

With multiple processors, you can also set processor affinity or *asymmetric multiprocessing* so that a processor is dedicated to a specific task of your choosing. For example, you could dedicate one processor to performing a complex scientific calculation while the other processor performs other server functions.

Hard Disk

Workstations and servers both store programs and data on hard disks. However, server hard disks are often optimized for high performance and **throughput**—a measure of the quantity of data sent or received in a second. For example, a 100 MBps throughput sends or receives at a maximum of 100 MB per second. In most computers, the hard disk is the slowest component and, therefore, a bottleneck that inhibits the overall performance of the system. Typically, the processor and memory can function thousands of times more quickly than the hard disk and often must wait for the hard disk to read or

write data before performing other tasks. The following characteristics contribute to improving hard disk performance:

- Servers are capable of implementing multiple hard disks while addressing them as a single logical disk. For example, drive E might appear to the NOS and users as a single drive. However, an administrator can use special hardware or software to implement several disks at once that act as a single drive E, thus reducing or eliminating a hard disk bottleneck. As a rough approximation, you could utilize two hard disks to double the performance of the logical disk. (Multiple use of disks in this manner is known as "striping" and is addressed in more detail in Chapter 5.)
- A **buffer** (or read cache) on a hard disk is memory that functions similarly to cache on a processor. When the hard disk retrieves data, it can store part of the data in a buffer. Later, the CPU can request the same data and it will be retrieved from the buffer, which is many times faster than mechanically retrieving it from the hard disk. Also, because data tends to be stored sequentially on the disk, a buffer can automatically read the next sequential data on the hard disk in anticipation that it might be needed soon after. An optimized workstation hard disk might have between 512 KB and 2 MB of buffer, while a hard disk optimized for server use might have 4 MB.
- **Access time** (also called "seek time") is the time it takes for the hard disk drive head, which reads or writes data on the platter, to arrive at the location of the data. Access time depends upon the spin rate of the hard disk. Most workstation hard disks spin at 5400 or 7200 revolutions per minute (rpm). Server hard disks usually spin at 7200, 10,000, or 15,000 rpm or higher. The faster the disk spins, the more quickly the required data can arrive under the disk drive head. Access times for lower-end workstation hard disks are around 9 or 10 milliseconds (ms), and high-end server hard disks can be as low as 4.5 ms.
- The disk drive **interface** is the hardware connecting the drive to the computer motherboard. The interface is perhaps the most important factor in hard disk performance. At the low end, the workstation interface is usually an Advanced Technology Attachment (ATA), considered synonymous with the Enhanced IDE (EIDE) interface, of which there are several types. Generally, workstation interfaces can read data from the hard disk at between 33 and 100 MBps and can connect up to four hard disks (two disks per channel). Server interfaces such as Small Computer Systems Interface (SCSI) can read data at around 100 MBps and support up to 15 hard disks.

Note

ATA interfaces are making great strides in performance such that they can be utilized in some servers to provide very good performance at an exceptionally low cost compared to high-end SCSI controllers.

In addition to the performance factors that distinguish server processors and hard disks, storage capacity is significantly greater on servers than on workstations.

Storage Capabilities

Network servers require more hard disk storage space than typical workstations in order to serve files and applications. A file server must have enough disk space available to store files for all users. Because most workstations have less storage space than a server, users could fill up their available local hard disk space over time. By storing files on the file server, an administrator can more feasibly monitor available hard disk space than he or she could by monitoring disk space on each workstation. Should low disk space occur on the server, the administrator could use any of several methods to add hard disks on the fly without any disruption in service. This is made possible by **hot-plug** or **hot-swap** capability, which means that you can add or remove a device without first powering down the computer.

Note

Software programs are available that can monitor each user workstation's disk space individually and notify administrators of low disk space. In addition, these programs (as well as Windows 2000) can restrict the amount of disk space available for each user on both the local workstation and the file server. However, it is better to minimize administration by using the same software to monitor a few file servers.

Access to Applications

Network applications are server-based programs that run in memory and on the processor on behalf of other servers or clients. Running applications on the server (referred to as **back-end applications**) minimizes or eliminates the processing, storage, and memory requirements for each individual client. Instead, the client computer runs a front-end application requiring much less overhead. (A **front-end application** is an application running on the client that retrieves information processed by a back-end application.) For example, it is much more effective to run most larger databases on a high-powered server than on client workstations. Even with multiple users running the database and performing queries at once, the server can usually outperform the same tasks on a workstation, because each workstation does not have to download the entire database across the network.

Let's say an organization has a relatively small database of 20 MB stored on a file server. Without using an application server, users at each workstation access the database file and run the database application locally. This means that each user accessing the database must download the entire 20 MB database, as illustrated in Figure 1-9. This is quite taxing on network bandwidth and requires users to wait for the download to complete. However, with an application server, the back-end database application and file stay on the server. If each record is 1 KB in size, and a user runs a query against the database with 30 matches, the client downloads only a 30 KB result as requested by the client's front-end database application (Figure 1-10).

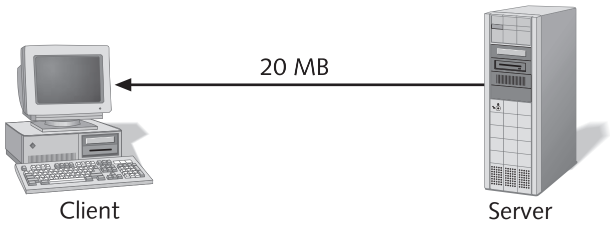

Figure 1-9 The entire database downloads to the client

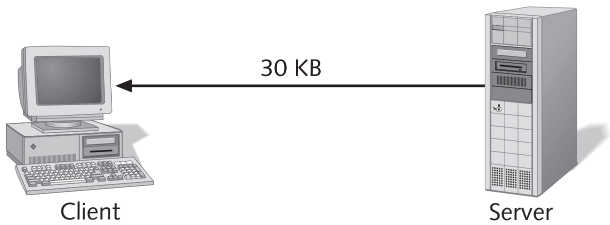

Figure 1-10 Only the database query results download to the client

There are several reasons for the performance gains with an application server, including the use of multiple processors, increased RAM, and fast hard disks on the server. Also, downloading only query results instead of the entire database greatly reduces the burden on network bandwidth. Examples of network applications and services include:

- *Database applications:* Programs such as Microsoft SQL 7 or SQL 2000 perform many of the functions and provide the benefits mentioned in the previous paragraph. Primarily, application databases benefit from the server in its role as a data warehouse (a storage location for extremely large databases) and from the server's processing power.
- *Email services:* Users can, to some degree, utilize Internet email with little if any involvement from the network administrator. However, an email server allows the administrator to perform management functions such as monitoring, virus scanning, forwarding, integration with directory services, clustering, security management, **failover** (an alternate system that takes over for a failed system), redundancy, and so forth. Many email servers go beyond email services to provide collaboration with users' calendars and videoconferencing.
- *Network management software:* This varies greatly from one product to another. Administrators of smaller networks might not require network management software, but larger networks require some sort of management software to manage client software distribution and licensing, monitor user activity, monitor and manage network traffic, back up and restore data, manage Internet services, integrate with other server operating systems, and more. Network management software is particularly useful because it can email or page the administrator when there is a problem with a server. Examples of network

management software include Microsoft SMS, Computer Associates Unicenter TNG, and IBM Tivoli.

- *Remote access:* Even when away from their desk or office, users require access to the LAN. When a user is away from the local network, he or she is a **remote user**. Many remote users telecommute from home or connect to the LAN from another office (Figure 1-11). Remote access makes this possible in two primary ways. First, the user can use a modem to dial into the network. The remote access server also has a modem dedicated to remote access purposes, and receives the user's call. Once the connection is complete, the user experiences the same network access as when locally present and connected to the LAN (except that network connectivity is only as fast as the modem, usually between 14.4 and 53 Kbps depending on phone line conditions). Second, faster remote access connections such as digital subscriber line (DSL) are replacing the modem-to-modem connection from client to server because they are faster and always on. However, traveling users generally continue to use a modem to connect (from a hotel room, for example).

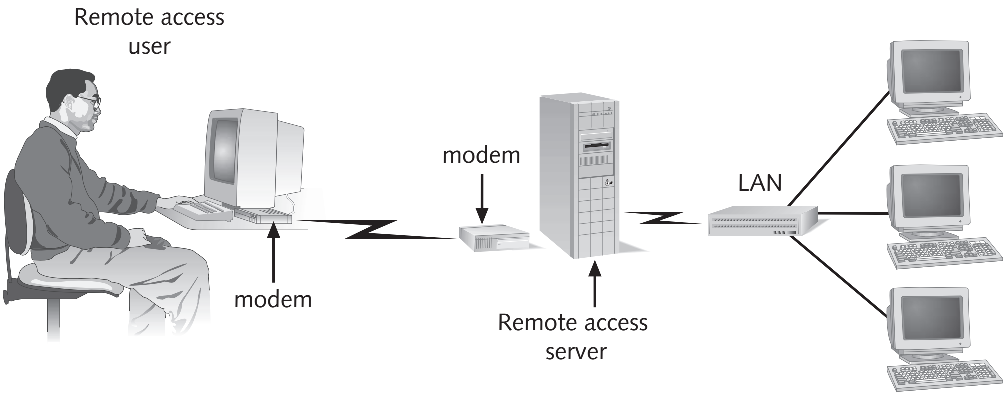

Figure 1-11 Remote users have the same network access as local users

- *Virtual private networks:* Users can also connect to the network over a **virtual private network (VPN)**, which is a highly secured network connection that makes eavesdropping from unauthorized persons nearly impossible. This connection can take place over the first connection type, a modem. Increasingly, however, users utilize an existing public Internet connection to establish the VPN connection. For example, a telecommuter with a high-speed Internet connection to his or her home office can establish a VPN connection to the corporate LAN and experience virtually the same security as if locally connected to the LAN. Windows 2000 utilizes routing and Remote Access Server (RAS) and Novell NetWare utilizes Network Access Server (NAS) to provide remote access.

Centralized Management

Any time the network grows beyond two or three dozen users, the administrator's role can become more reactive than proactive. That's because he or she must continually go from one computer to the next, troubleshooting applications, helping users find their files, and so forth. Although adding a server does not eliminate these tasks, its centralized management capabilities can certainly contribute to reducing overall administration.

Centralized management means that administrators can administer servers from one place instead of traversing frantically from one end of the building or campus to the other fixing problems. With the proper NOS options enabled or network management software installed on a server, an administrator can browse a user's hard disk from the administrator's own computer to help them find the files they are looking for. Or if a user is experiencing difficulty using a program, the administrator can view the user's desktop without leaving his or her own seat. Sharing applications from the server also reduces administration because the administrator can install, troubleshoot, and upgrade applications from a single server location rather than on each individual client.

Caution

Managing applications on a single server reduces administration, but it also provides a single point of failure. If the server fails or the application stops functioning correctly, all network users are affected. Consider having a redundancy plan. For example, multiple servers hosting the application can allow users to continue using the application should one of the servers fail (Figure 1-12). Clustering is an example of **fault tolerance**, which allows for continued service despite failure of a server or component.

Cost Effectiveness

While you, as the network administrator, see obvious benefits for installing a server, it might not be up to you. Much of the time, some combination of management, accountants, or both determine what they think is best for the network strictly on the basis of how much something costs. So when you propose adding a modest $8,000 to the budget for a server and applications for your LAN of 30 users, be prepared for the possibility of a knee-jerk reaction and a resounding no.

Increasingly, administrators must show the money sense in the budgets they propose, and in fairness to an organization's financial affairs, they should. As you prepare to propose new server equipment, know that well-planned servers eventually recoup their initial expense and continued operation in long-term cost savings and increased productivity. The following sections illustrate the savings in several important areas.

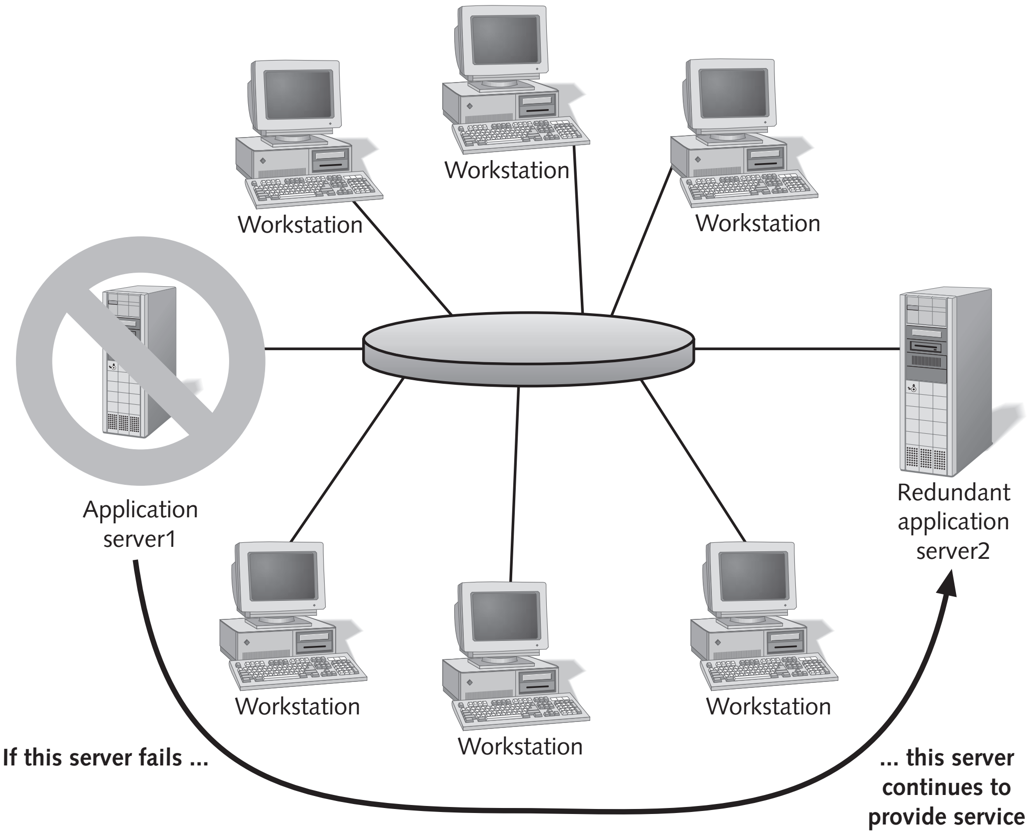

Figure 1-12 Clustering ensures continued service to clients

Storage

Workstations in many organizations are several years old. Applications, operating systems, and data files on the workstations grow larger, but the hard disks do not. It would be much more cost effective to install centralized storage on a server for all the users. Although the price per megabyte on high-performance server hard disks might be higher than that of a workstation, you can save time (and hence money) by performing a single hard disk upgrade on the server instead of upgrading each workstation's hard disk. Upgrading individual user's hard disks might also involve the sometimes tricky process of moving the existing operating system, applications, and data to the new drive without disruption or loss of current settings.

TIP

Sometimes you cannot avoid transferring the contents of the user's current hard disk to a new one. The best way to perform this operation safely is to use software designed for this purpose. Several products from third parties allow you to do this. For example, PowerQuest's Drive Image Pro 4.0 can duplicate one hard disk to another.

Processors and Memory

As operating systems and applications become increasingly complex, workstations require more processing power and RAM. Individual workstation upgrades quickly become cost prohibitive. However, you can upgrade one or more servers and transfer the memory and processing burden to the server, as in the case of a large database. Also, to avoid the costly operating system and hardware upgrade for client computers, you can leave the existing operating system and hardware, and make the workstations thin clients. A **thin client** (similar to a dumb terminal) receives its operating system environment, including applications and data, from the server (Figure 1-13). The server does all the work, and the client sends the input from the keyboard and monitor and receives the output to the monitor. For example, users could boot to Microsoft Windows 3.11 at the workstation and connect to the server to run the Windows 2000 operating system as the primary environment. Microsoft Windows NT Server 4.0, Terminal Server Edition, Microsoft Windows 2000 Server, and Citrix MetaFrame provide thin client services. Novell's ZENworks integrates with and enhances thin client services to provide powerful administrative tools and to better manage the user desktop.

Be careful: If you plan to move to a thin client solution, the network utilization will increase significantly. You might have to anticipate this and increase available bandwidth prior to moving to thin clients.

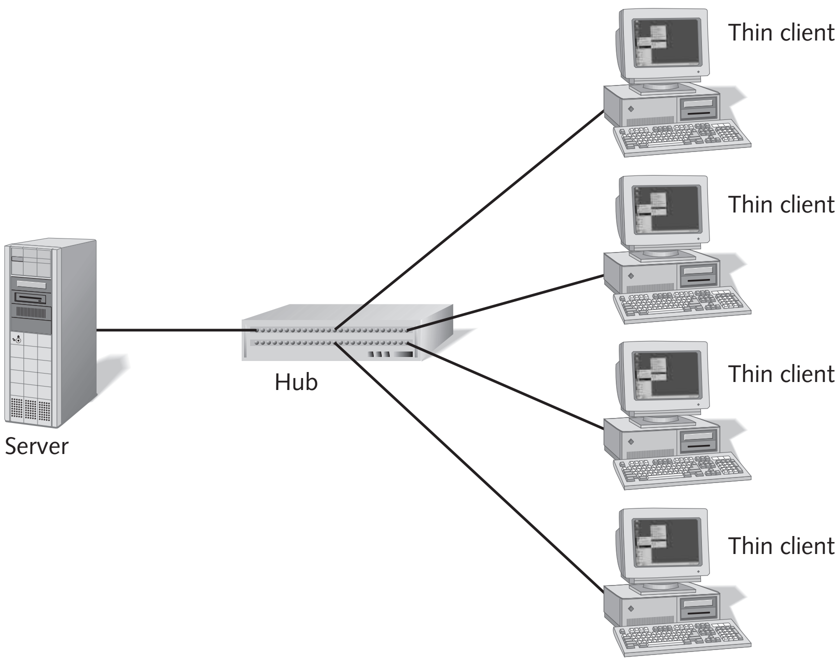

Figure 1-13 Each thin client receives its desktop environment from the server

Savings, Efficiency, and Productivity

Efficiency and productivity are difficult to quantify and vary from one organization to the next. Most managers and accountants to whom administrators must defend their budgets already understand that computers are generally cost effective. When proposing to add server equipment or software, you must demonstrate that the addition saves money, increases efficiency and productivity, or both.

For example, what would it cost a corporation to lose its customer database to a competitor? A company with an Internet connection exposes itself to unauthorized access to the world at large. Intruders who succeed in invading your network might not do anything harmful, but administrators must plan for the worst in case someone destroys, steals, or alters data. An administrator can add a **firewall**, which protects the internal LAN from the public Internet and is placed between the LAN and the Internet (Figure 1-14). By protecting the network, a firewall saves the company from potential financial and productivity losses. A firewall can be either a hardware or software solution. In a server context, it is a server with special software. For example, Novell offers BorderManager as its software firewall.

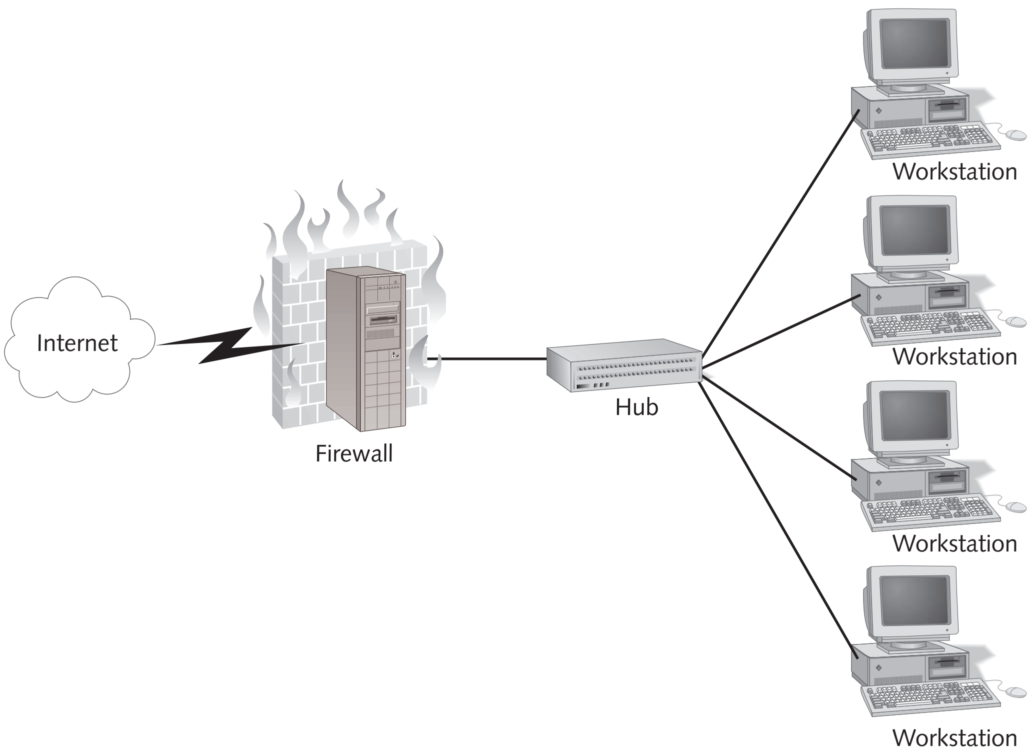

Figure 1-14 A firewall protects user workstations from the Internet

Classes of Servers

Having established the need for a server in several contexts, you must also determine what type of server is appropriate for your organization. A poorly chosen server might be more expensive and more powerful than current and future needs require, or it might be woefully inadequate, necessitating the purchase of more equipment at a less efficient cost factor. For example, if an administrator purchased a single PC server capable of supporting two processors, and after installation discovers that it cannot meet the processing demands placed upon it, he or she must purchase another system. In this case, it might have been less expensive to purchase a single midrange server with support for four or more processors in the first place.

Although the Server+ certification focuses primarily on PC servers, a brief discussion of the mainframe, minicomputer, and midrange computer classes provides you with a broader understanding of other server types. In addition, it helps to know something about mainframes because administrators often connect PC servers to them.

PC Servers

Basically, a PC server uses Intel-compatible complex instruction set computing (CISC) processors, as opposed to reduced instruction set computing (RISC) processors from vendors such as Sun Microsystems or IBM. At the low end, a PC server can be a workstation-level computer in terms of its hardware. For example, a small insurance company with 12 desktop PC workstations has a need for a file and logon server. The company could install a NOS on one of the existing PCs and suddenly the PC becomes a server, and the network changes from a peer-to-peer LAN to a client-server network. The fact that the "server" has only desktop PC-level hardware does not change the fact that it is a server, and in this small LAN, more powerful hardware is unnecessary.

On the high end, a PC server can involve significantly more powerful hardware; therefore, it has a much larger case than a standard PC workstation. (For purposes of this book, assume that the server is a high-end server unless stated otherwise.) The server case also provides much more space for expansion and requires more components, such as extra power supplies.

Other characteristics of a high-end server include:

- *Hot-pluggable PCI slots:* Peripheral component interface (PCI) slots store various adapters such as disk controllers, video cards, and network adapters. A desktop PC workstation would require you to shut down the computer before adding or removing adapters. A server with hot-pluggable PCI slots allows you to add or remove adapters while the server is on. This is a very useful feature in mission-critical operations where you cannot afford the extra minutes required to shut down and restart the server. Some servers offer this functionality by allowing you to shut down individual PCI slots without powering down the entire system.

- *Hot-swappable hard disks:* These allow you to quickly remove hard disks without powering down the system (see Figure 1-15). They are useful for quick emergencies that occur when it's time to upgrade a hard disk or replace a failed one.
- *Power supply:* Servers usually have two or three power supplies to power the additional component slots and drives, and to provide redundancy. If one power supply fails, the system can usually continue to operate until you replace it with another hot-swappable power supply (Figure 1-16).

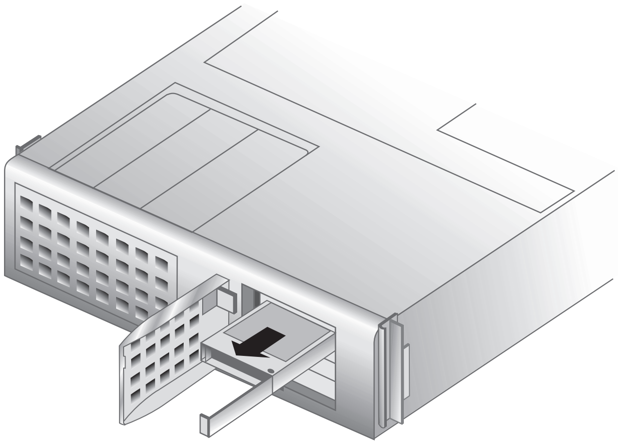

Figure 1-15 Hot-swapping hard disks

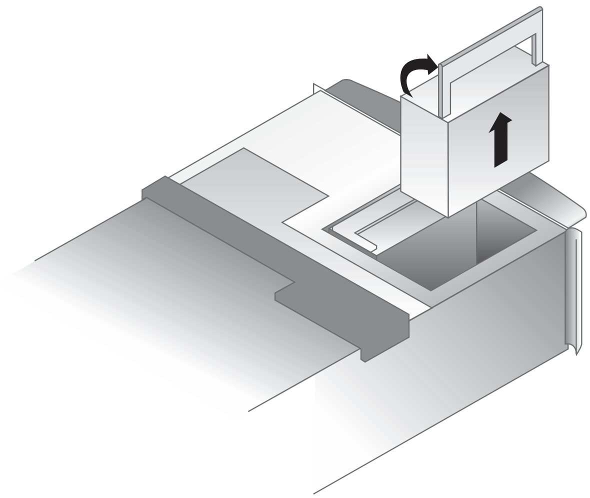

Figure 1-16 Hot-swapping a power supply

- *Cooling fans:* Redundant hot-swappable cooling fans help ensure that the system does not overheat if one of the fans fails, because the other fan(s) can continue operating until you hot-swap the failed fan for a new one. This is a critical element to maintaining server uptime. (**Uptime** is the continued operation of the overall server or specific components such as the hard disk, depending upon the context.) An overheated system can quickly burn out or reduce the life of the components, especially the processor, which is the hottest component of the system.
- *Expandability:* Servers have more PCI slots and drive bays, and are modularly designed so that you can easily add or remove the components (Figure 1-17).

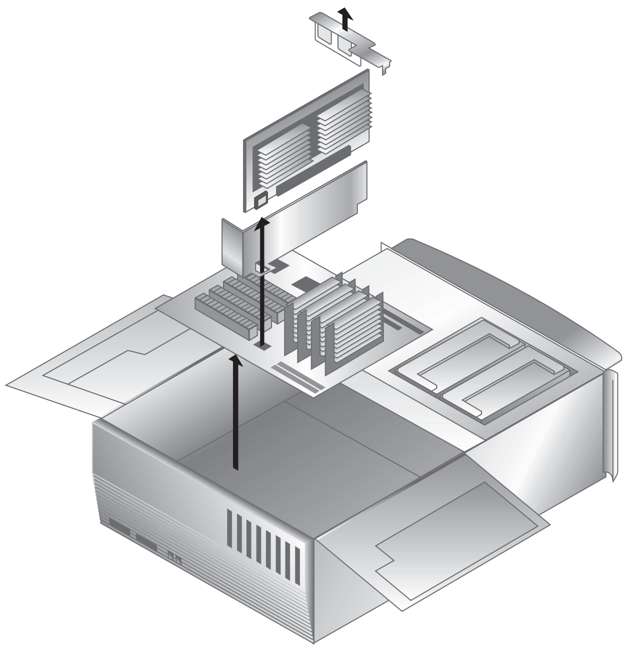

Figure 1-17 Servers are modular to facilitate adding and removing parts

- *Heavy-duty chassis:* The **chassis** is the metal frame to which the motherboard is attached and which forms the case structure. A server chassis is much heavier than a workstation chassis and often weighs about 75 lbs, including the case and components. It is heavier because the metal is of a heavier gauge, and the case is often larger to accommodate more internal components, such as more hard disks. After adding all internal equipment (drives, power supplies, and so on), the server can weigh over 170 lbs!
- *Rack mountable:* For many servers, when you order the server you can specify a **tower** (upright, free-standing case) or a **rack** configuration. The rack mount consolidates space because you can stack several servers and other network equipment in the same floor space, offering good **density** (Figure 1-18).

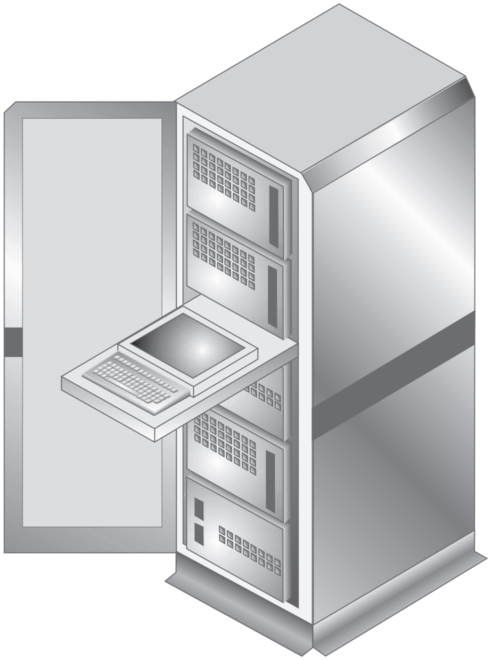

Figure 1-18 A rack stacks servers and equipment in a small floor space

Mainframe or "Big Iron"

A **mainframe** is a large and extremely powerful computer. Nearly all computers in the 1960s were mainframes and usually filled an entire room. IBM was one of the first mainframe pioneers and continues to lead the mainframe market today with its S/390 and ZSeries 900 systems. Mainframe computers are colloquially known as "big iron" in reference to their size and power. The term "mainframe" becomes less definitive as minicomputers become increasingly powerful and some mainframes become smaller in size. However, there are several characteristics that continue to distinguish a mainframe:

- *Size:* It's just plain big. Many of the newest mainframes boast a relatively small size compared to their predecessors, and require about 20 square feet of floor space, and weigh over 2000 lbs.
- *Cost:* Mainframe prices are in the range of "if you have to ask how much, you can't afford it." Mainframes are still in the range of at least several hundred thousand dollars.
- *Processors:* Mainframes nearly always use multiple processors—up to 32 for Intel-based machines. Also, mainframes frequently use proprietary RISC-based processors that usually run a version of UNIX. Although the processors in a mainframe might be comparable to a PC in terms of MHz, a PC processor is busy delivering data to and from its peripherals. This is very time consuming

in terms of clock cycles. The mainframe utilizes mechanisms to offload I/O and avoid the negative impact on performance.

- *Bus speed:* No computer can be any faster than its main board (or motherboard) bus speed. While a PC server might have a speed of 100 MHz, a mainframe might have significantly higher bus speed, which when multiplied with its SMP capability can provide blazing processing performance.
- *Memory:* Mainframes can have upward of 96 GB of main memory. Also, mainframes have extremely well-engineered cache and memory architecture, providing very high efficiency and performance.
- *Storage:* Mainframe disk storage is capable of storing hundreds of GB or a terabyte (TB) or more. Typically, large quantities of storage are external to the server.
- *Durability:* Mainframe hardware circuitry design detects and corrects errors, and it can anticipate and alert administrators of impending problems. Some mainframes can generate a list of parts that need replacement at the next regular maintenance. **Mean time between failure (MTBF)** is the anticipated lifetime of a computer (or one of its components). Mainframes often have MTBF of about 20 years. While it is difficult to specify PC server MTBF, it is usually a fraction of that of a mainframe.

The Minicomputer or Midrange Computer

Minicomputer characteristics lie somewhere between the desktop workstation and the mainframe computer. The minicomputer is sometimes known as the **midrange computer**. This is an extremely broad range, with distinction between midrange and PC blurring at the low end, and between midrange and mainframe at the high end. The first minicomputer was the Digital Equipment Corporation (DEC) PDP-1, which cost $20,000 in 1959 and was not as large as a mainframe but was still quite sizable. Compared to mainframes costing millions of dollars, the price tag was extremely attractive. However, the PDP series was not considered successful until the mass-produced PDP-8 in 1965, which was a tabletop-size unit. Minicomputers also offer an advantage compared to mainframes in that they require significantly less floor space and weigh less.

This book and the CompTIA Server+ Certification focus on PC servers.

Chapter Summary

- A peer-to-peer network is a collection of networked computers with no logon server to verify the identity of users. This model is called a peer-to-peer network because each network device has an equal (peer) level of authority. Peer-to-peer networks have limited growth potential because client workstations are limited to 10 concurrent inbound network connections. Peer-to-peer networks have weaknesses in regard to

security and file version consistency, and there is no centralized location on which to store user accounts.

- The client-server network begins with a LAN and one or more servers. The client-server network can also encompass a more complicated network configuration such as multiple, geographically distant LANs connected to one another across a relatively great distance, known as a wide area network (WAN).
- A server possesses more processing power, RAM, and hard disk capacity than workstation computers on a LAN. The server also has a server network operating system (NOS) such as Microsoft Windows NT or 2000, Linux, IBM OS/2, or Novell NetWare. A server running a NOS provides file and printer sharing, centralized file storage, security, and other network services.
- Three primary factors—clock speed, data bus, and cache—contribute to the effective speed (performance) of the processor or CPU. Servers generally have higher-performance CPUs than workstations, which means they have greater clock speed, bandwidth, and cache. Servers also may utilize symmetric multiprocessing (SMP), the simultaneous use of multiple processors on the same server.
- Servers are capable of implementing multiple hard disks while addressing them as a single logical disk, which reduces or eliminates hard disk bottlenecks. Server hard disks also utilize buffers and interfaces that can easily support multiple hard disks. By storing files on the file server, an administrator can more feasibly monitor available hard disk space than he or she could by monitoring disk space on each workstation.
- Network applications are server-based programs that run in memory and on the processor on behalf of other servers or clients. Running the application on the server minimizes or eliminates the processing, storage, and memory requirements for each client.
- Centralized management consolidates administration of user workstations and other servers to a single location.
- Increasingly, administrators must justify the budgets they propose. Well-planned servers eventually recoup their initial expense and continued operation in long-term cost savings and increased productivity.
- Different classes of servers include PC servers, mainframes, and minicomputer or midrange servers. A PC server generally uses Intel-compatible complex instruction set computing (CISC) processors as opposed to reduced instruction set computing (RISC) processors from vendors such as Sun or IBM. Characteristics of a high-end server include hot-pluggable PCI slots, hot-swappable hard disks, redundant power supplies, redundant cooling fans, high expandability, a heavy-duty chassis, and rack mount capability.

Key Terms

access time — The time it takes for the hard disk drive head to arrive at the location of the data. Access time depends upon the spin rate of the hard disk.

authentication — Verification of a person's identity based on their credentials (usually a username and password).

authorization — Verification that an authenticated user is permitted to access a network resource.

back-end application — Applications that run on the server on behalf of the client.

buffer (or **read cache**) — On a hard disk, memory that stores part of the data read from the hard disk. Later, the CPU can request the same data and it will be retrieved from the buffer, which is many times faster than mechanically retrieving it from the hard disk.

cache — Memory that assists performance by storing frequently used data for fast access. Processors and hard disks use cache, and cache is separate from main system memory (RAM).

centralized management — The ability to administer a given system from a single location instead of disparate locations.

chassis — The metal frame to which the motherboard is attached and which forms the case structure.

client — A network workstation that requests and receives service from the server.

client-server — A network that begins with a LAN and one or more servers. The client-server network can also encompass a more complicated network configuration such as multiple, geographically distant LANs connected to one another across a relatively great distance known as a wide area network (WAN).

clock speed — The number of instructions the processor can execute in a single instruction, measured in megahertz (MHz)—which is one million cycles per second. Instructions sent to the processor require a certain number of cycles, so the more cycles the processor can handle per second, the faster it operates, or "thinks."

clustering — Redundant servers hosting the same application for the purpose of fault tolerance. If one of the servers fails, the remaining server(s) continues to serve the application to the network.

data bus — The number of bits the processor can execute in a single instruction. Bandwidth is typically 32 bits. Some new processors offer 64-bit bandwidth.

density — A term used with equipment racks that describes consolidation of space because you can stack several servers and other network equipment into the rack in the same floor space.

failover — An alternate system that takes over for a failed system.

fault tolerance — Continued service despite failure of a server or component.

file server — A server that provides a central location to store files for network clients.

firewall — A hardware or software solution that protects internal LAN users from the public Internet.

front-end application — An application running on the client that retrieves information processed by a back-end application.

host — A network device (usually a computer) in a TCP/IP network.

hot-plug (or **hot-swap)** — Add or remove a device without first powering down the computer.

hub — A network device that connects network cables together in a central, star configuration. Passive hubs simply make the connections, and active hubs (multiport repeaters) regenerate the signal to increase the distance it can travel.

interface — The hardware connecting the drive to the computer motherboard.

local area network (LAN) — A collection of computers in close proximity to one another on a single network.

mainframe — The most powerful level of computer classification, mainframes are extremely large and powerful computers. Also known as "big iron."

mean time between failure (MTBF) — The anticipated lifetime of a computer or one of its components.

midrange computer (or **minicomputer)** — A broad computer classification that lies somewhere between desktop workstation and mainframe computer.

network — A collection of two or more computers connected with transmission media such as network cable or wireless means, such as radio or infrared signals. Usually includes other devices such as printers.

network applications — Server-based programs that run in memory and on the processor on behalf of other servers or clients.

network device — Any device connected to the network for purposes of communicating with other network devices. (A network device is also known as a host in most networks.)

network operating system (NOS) — Provides file and printer sharing, centralized file storage, security, and various services. Primary examples of a NOS include Microsoft Windows NT or 2000, Linux, IBM OS/2, or Novell NetWare.

network resource — An object users can access from across the network, such as printers, files, and folders.

peer-to-peer network — A collection of networked computers with no logon server to verify the identity of users. Each network device has an equal (peer) level of authority.

permissions — The configured level of access applied to a resource. For example, if a user can read a file but not change the file, then they have read-only permission.

rack — A cabinet that houses stacked network equipment, storage, and servers. A rack can store multiple items in the same floor space.

redundancy — The ability to continue providing service when something fails. For example, if a hard disk fails, a redundant hard disk can continue to store and serve files.

remote user — A user connected to the LAN from a geographically distant location, usually over a modem or virtual private network (VPN).

server — A computer with more processing power, RAM, and hard disk capacity than typical workstations. A server has a server NOS such as Microsoft Windows NT or Novell NetWare, and provides file and printer sharing, centralized file storage, security, and various services.

services — A function of the NOS that provides server features to the network.

superscalar — A processor architecture that allows a processor to execute more than one instruction in a single clock cycle.

switch — Similar to a router in that it segments a network, and similar to a hub in that it connects network cables together in a central, star configuration. Switches forward traffic at very high speeds.

symmetric multiprocessing (SMP) — The simultaneous use of multiple processors on the same server.

thin client — A computer that receives its operating system environment, including applications and data, from the server.

throughput — A measure of the quantity of data sent or received in a second.

tower — An upright, free-standing computer case.

uptime — The continued operation of the overall server or specific components.

virtual private network (VPN) — A highly secured network connection over an otherwise unsecured network such as the Internet.

wide area network (WAN) — Multiple, geographically distant LANs connected to one another across a relatively great distance.

workstation — Desktop computer with only enough hardware to service the needs of a single user at a time. Synonymous in most contexts with PC, desktop computer, or client.

Review Questions

1. A network is:
   a. at least a dozen computers on one network cable
   b. any context in which the server and client can communicate with one another
   c. a collection of two or more computers connected with transmission media
   d. two or more computers that are in close proximity to one another
2. Which of the following are true of LANs and WANs? (Choose all that apply.)
   a. A LAN involves networked computers in close proximity to one another.
   b. A WAN involves at least one Internet connection.
   c. A LAN is a linked area network involving two or more sites.
   d. A WAN involves two or more geographically separate LANs connected to one another.

3. Which of the following is true of a peer-to-peer network?
   a. It is limited to no more than one logon server.
   b. It is limited to no more than one file server.
   c. No logon server is present.
   d. No server of any kind is present.
4. Why not use Microsoft Windows NT or Windows 2000 workstations to create user accounts instead of a server?
   a. There is no support for complex passwords.
   b. There is no support for encrypted passwords.
   c. More administration is necessary. Each workstation must match the username and password exactly for users to gain access to resources.
   d. Each workstation must copy the username and password to the other workstations.
5. Which of the following is not a server NOS?
   a. Linux
   b. IBM OS/2
   c. Windows Millennium Edition
   d. Novell NetWare
6. Which of the following is not a feature of a NOS?
   a. email
   b. resource sharing
   c. security
   d. multiple simultaneous connections
7. How does a smart card increase network security?
   a. It provides a better online shopping experience.
   b. You cannot enter the building without it.
   c. The user must possess the smart card to log on.
   d. Users do not need to use a password.
8. What is the difference between authentication and authorization?
   a. Authentication verifies the identity of the user and authorization permits access to resources.
   b. Authorization verifies the identity of the user and authentication permits access to resources.
   c. Authentication validates computers and authorization validates users.
   d. Authentication and authorization are the same thing.

9. You recently discovered that some users are changing their own salaries by accessing the Salaries.xls file. What should you do?
   a. Change the user's password.
   b. Change the permissions on the Salaries.xls file.
   c. Deny users permission to log on to the file server.
   d. Fire the users.
10. Why wouldn't a wise administrator allow users to store data files on their own workstations? (Choose all that apply.)
    a. Users might not assign prudent permissions to their files.
    b. Backing up the files is more difficult.
    c. Workstations might require a hard disk upgrade to store the files.
    d. A workstation operating system might not permit sufficient simultaneous network connections for shared files.
11. Which of the following does not contribute to the effective speed of the processor?
    a. clock speed
    b. cache
    c. voltage of the power supply
    d. data bus
12. Hard disk and processor cache provides which of the following benefits?
    a. fast access to recently accessed data
    b. long-term storage
    c. on-the-fly data backup
    d. increase in main system RAM
13. How does SMP benefit the performance of a server?
    a. Multiple hard disks aggregate read/write performance.
    b. Additional cache improves data access.
    c. Multiple processors increase processor performance.
    d. Multiple computers provide clustering in case of server failure.
14. The most likely bottleneck in most computers is:
    a. the processor
    b. the hard disk
    c. the main memory
    d. the cache memory

15. Access time is dependent upon:
    a. how fast the processor can issue instructions to the hard disk
    b. the number of platters the hard disk uses
    c. the disk drive interface
    d. how fast the hard disk spins
16. Several server components offer hot-swap capability, which is:
    a. the ability to replace a hardware server in the rack in less than five minutes
    b. a temperature problem that requires immediate replacement of a processor when it overheats
    c. the ability to distribute processing tasks between two processors
    d. the ability to replace a hardware component without first turning off server power
17. Which one of the following statements is true of an application server?
    a. Applications run on the clients, freeing the server resources to perform more complex tasks.
    b. Applications run on the server, improving responsiveness for the clients.
    c. Applications run on the server, increasing network traffic.
    d. Application installation files are located on the server and installed on the client.
18. Which of the following provides thin client services?
    a. Novell Network Access Server
    b. Windows NT 4.0 or Windows 2000 remote access server
    c. Citrix MetaFrame
    d. any workstation operating system
19. Which of the following characterize a PC server?
    a. CISC processors
    b. RISC processors
    c. extremely heavy
    d. utilizes a FEP
20. Which of the following characterize a mainframe?
    a. the largest classification in terms of size
    b. hard disk storage up to hundreds of GB
    c. utilize a single extremely powerful RISC-based processor
    d. usually cost less than $100,000

Hands-on Projects

Each chapter of this book contains hands-on projects in which you can verify and practice concepts, techniques, and other information discussed in the main text. These projects are intended to give you direct experience and to reinforce your learning. Plan to create a lab journal or a running word-processed document so that you can record your findings as you perform each project. The lab journal or word-processed document will be a valuable study aid.

Project 1-1

Most desktop operating systems such as Windows 9x (referring to Windows 95, 98, or ME) allow only a limited number of simultaneous connections to a shared resource. This can be extremely limiting in networks of a dozen or more users. In this project, you will create a share and observe the connection limitations of a peer-to-peer operating system.

1. From the Windows 9x computer, double-click **My Computer**.
2. Double-click the **C:** drive.
3. Right-click in a blank location of the window, and from the context menu, click **New**, and then click **Folder**.
4. Type ***YourName*Share** as the name of the folder.
5. Right-click your new folder and click **Sharing** from the context menu.

> The Sharing option is only available in the context menu if the "allow others access to my files" option has been enabled in File and Printer Sharing.

6. In the Sharing dialog box, click the **Shared As** radio button.
7. Click **OK**.
8. Close all open windows.
9. Double-click **Network Neighborhood**, and browse other computers. Notice that others now have shared folders that you can access.
10. Double-click a specific computer as determined by your instructor. If you have 11 or more students, the eleventh simultaneous connection will fail to connect.
11. Close all open windows.

Project 1-2

A server does not offer the same limitations as a workstation desktop operating system. In this project, your instructor will add a server. Then, you will access the server and add a file. In contrast to the peer-to-peer model, the client-server model allows multiple simultaneous connections.

1. The instructor adds a server to the network.

2. The instructor creates a share called ClassShare.
3. As per the instructor's direction, double-click **Network Neighborhood** to locate the server, and locate the share.
4. Notice that more than 10 users can connect to the server share.
5. Create a text file named ***YourName*** inside the shared folder, enter some text in the file, and then save and close the file.
6. Leave the Network Neighborhood windows open for Hands-on Project 1-3.
7. In your lab journal or word-processed file, list some of the benefits of accessing shared files from a server instead of a workstation.

Hands-on Project

Project 1-3

In this project, you will see how servers can adjust security settings by adjusting shared permissions. Your instructor will change permissions on the share you accessed earlier. Observe how this affects your level of access.

1. Double-click the file you created in Project 1-1.
2. Make changes to the file and attempt to save it. You cannot save the file because the instructor changed the permissions so that you can only read (not write to) the file.
3. Close all open windows and applications.

Hands-on Project

Project 1-4

Administering servers requires you to keep abreast of current servers and equipment. If an organization wants you to add a server, you must generally have a grasp of what is available and how much it might cost. If an organization wants you to improve its current network, you must know what kind of equipment you can add, and if you should add or replace servers.

Note

The web sites listed in these projects and throughout the book are usually the vendor's home pages, because specific web pages might be outdated by the time you read this. Go to the vendor's home page and follow general steps to arrive at the requested locations.

Use your web browser to visit *www.dell.com*. Look for a server for medium-size to large businesses and answer the following questions in your lab journal or word-processed document:

1. What appears to be the most powerful server Dell offers?
2. What is the maximum number of processors this server can use?
3. What components are hot-pluggable?
4. Is the server rack mount only or can you get it in both rack and tower configurations?

5. How much does it cost?
6. What are a few of the server's features that cause it to stand out?
7. Would you consider this server to be a PC server, a midrange server, or a mainframe?

Project 1-5

This time, you will visit a different web site and notice a difference in equipment.

Use your web browser to visit *www.unisys.com*. Look for the server content on the site, and answer the following questions:

1. What appears to be the most powerful server Unisys offers?
2. What is the maximum number of processors this server can use?
3. What components are hot-pluggable?
4. How much does the server weigh?
5. How much does it cost? (Prices might not be available because you usually must request a salesperson to bid a price for you.)
6. What are a few of the server's features that cause it to stand out?
7. Would you consider this server to be a PC server, a midrange server, or a mainframe?

Project 1-6

The focus of this book is to prepare students for the CompTIA Server+ exam. Although this book addresses the exam objectives, you should also view them at CompTIA's web site. These steps direct you to specific links. If the links have changed, attempt to locate the same general area.

1. Browse to *www.comptia.org.*
2. Click the **Certification** link. You see several certifications such as A+, Network+, and Server+.
3. Click the link for **Server+**. What kind of experience does CompTIA expect that candidates for the Server+ exam will have?
4. Notice that several vendors support the Server+ certification. Who are some of the vendors?
5. Some portions of the Server+ certification have a greater emphasis than others. Name three certification objectives that have the greatest emphasis.
6. Which other certifications does CompTIA recommend as prior experience for the Server+ exam?

Case Projects

1. Stan, the owner of a local vending machine supplier, calls upon your services to assist with their network of 14 desktop PCs. Stan has little, if any, knowledge about computers and even less about networks and servers, but he knows that all his PCs run Windows 98. Stan is concerned about security because he wants to allow Internet access to each desktop, but he is afraid of possible security risks. All his customers call in their orders when they need something, but Stan wants to eventually allow customers to use the web to place orders. Stan has limited financial resources to contribute to this, and does not expect to hire new employees or significantly expand his business over the next two or three years. Does Stan have a peer-to-peer or client-server network? What can you recommend to Stan to help him accomplish his objectives?
2. Andie, a stockbroker at a local stock brokerage, calls you to assist with their network. The brokerage is an office of about 100 users, each with workstations. A high-speed dedicated line connects the brokerage to the main headquarters in another city. The brokerage also has high-speed Internet access. The problem is, the computer that directly connects to the Internet connection failed a few days ago. The network administrator was on a sailing vacation and unreachable, so someone with little computing knowledge decided to plug the incoming line from the Internet directly into the company's rack of hubs. Nobody seemed to mind because all users were immediately productive and access to the Internet resumed. Now, Andie notices that several files on her Windows 98 computer are missing. She cannot find them in the Recycle Bin and is sure she did not delete them. Some of her client files have missing data, and others appear to have been tampered with. Other stockbrokers at the office are complaining of the same problems. What can you do to help? What is the most likely cause of these problems? What would have been the best course of action when the server connected to the Internet failed? What was the purpose of the server connected to the Internet?